Solution Overview

Accelerating AI Storage Access with the NVIDIA BlueField DPU

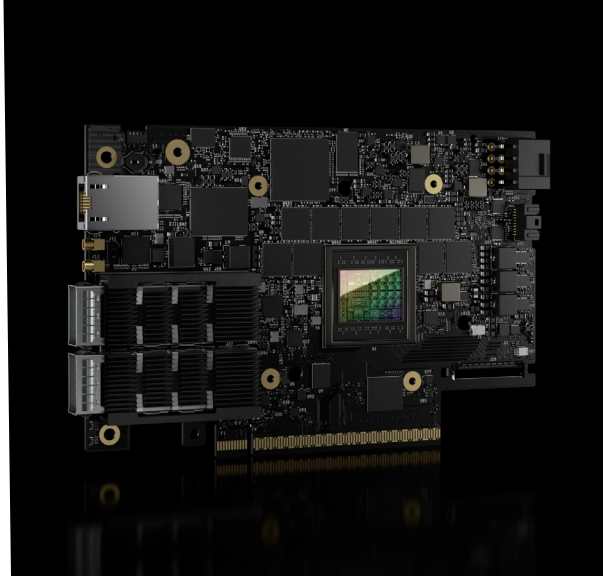

As data center storage becomes faster and more focused on supporting AI applications, the NVIDIA® BlueField® DPU is now the ideal infrastructure and networking accelerator for both storage arrays and the servers that connect to them.

Changes to the World of Storage

The world of data center storage has undergone three major changes in the last few years. First, the storage systems have completed a transition to flash, meaning all storage—except for some archival and backup storage—is 100 percent flash, supporting higher transaction (IOPS) and throughput rates. Second, storage systems now require faster network speeds of 100, 200, or even 400 Gigabits/second (Gb/s), to keep up with both the flash media and the faster servers. Third, an increasing number of data center storage deployments are dedicated solely to supporting AI—especially generative AI—which demands high performance from both storage and storage networking.

Following these pivotal changes in data center storage, NVIDIA BlueField DPUs emerge as the ideal solution for storage network traffic between AI platforms, standard servers, and storage.

Supporting AI Applications and AI Platforms

Modern AI applications, such as large language models (LLMs), are distributed across multiple GPU servers. The amount of data is frequently very large for training—dozens or hundreds of terabytes. Both training and interference may require data parallelism or model parallelism, meaning many GPU servers must access the same data. Training often requires multiple passes or epochs, where GPUs must reload the same data multiple times to refine the model. Training also requires frequent, sustained bursts of write I/O for checkpointing or logging. As a result, GPU nodes typically require high-throughput, low-latency storage access at speeds of 200Gb/s or 400 Gb/s, and these speeds are expected to increase in the future.

Many of the most popular AI systems, like NVIDIA DGX™ and many servers based on the NVIDIA HGX™ platform, already use BlueField DPUs for north-south network connectivity. They can also use NVIDIA® GPUDirect® Storage (GDS) to provide a direct data path between the GPU platforms and storage—without burdening the server CPUs with interrupts and data copies.

Challenges for High-Performance Storage

- **Performance:** Delivering high throughput and shared access to many servers simultaneously
- **Scalability:** Scaling performance across distributed compute and distributed storage clusters
- **Latency:** Maintaining consistently low latency across the network
- **Acceleration:** Supporting AI- and HPC-specific network accelerations

BlueField DPU Benefits for Storage

- Optimizes storage networking performance for AI
- Makes storage controllers more efficient and secure
- Simplifies storage provisioning and management for servers
- Already the top choice for many AI servers and storage solutions

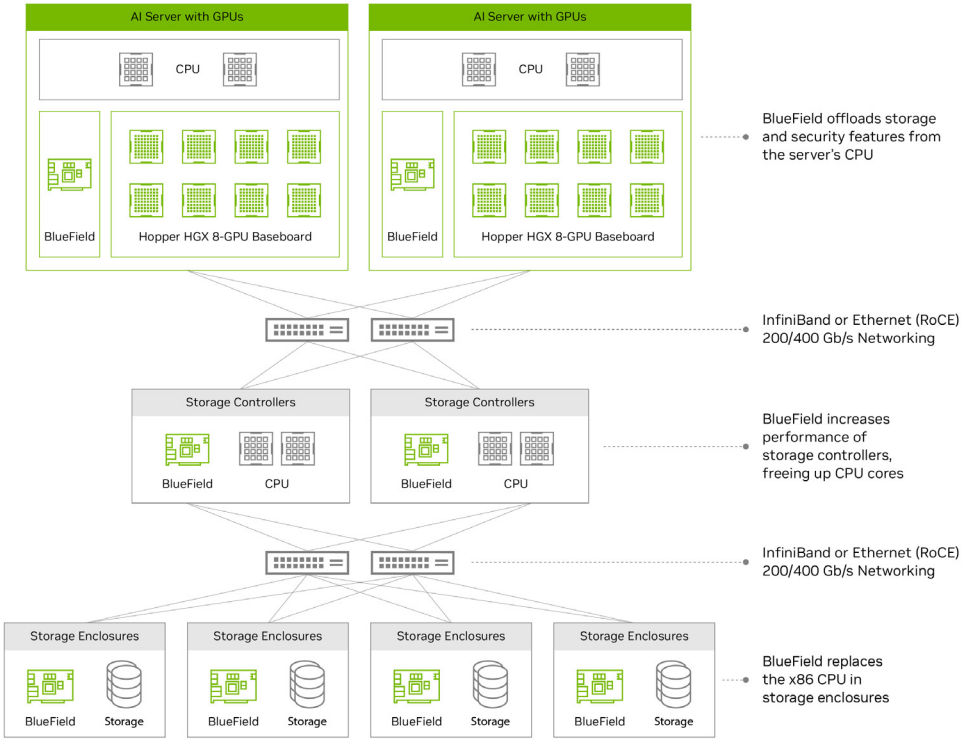

BlueField enhances storage performance, security, and efficiency in the server, storage controller, and storage enclosure.

Support for RoCE, InfiniBand, and the Fastest Network Speeds

Storage access should use InfiniBand or RDMA over Converged Ethernet (RoCE) networking for AI and HPC applications (required for GDS). AI storage benefits from RoCE extensions and MPI collective offloads, and large AI clusters demand the use of congestion control and adaptive routing for the storage network to minimize latency and maximize throughput. Only the BlueField DPUs offer full support for both InfiniBand and Ethernet connectivity at the highest available speeds—currently 400Gb/s and expected to increase soon. While many adapters have some support for RoCE, only BlueField DPUs deliver the most proven support for RoCE at the highest speeds, as well as support for features such as programmable congestion control, adaptive routing, and Zero-Touch RoCE.

Making Storage More Efficient

Storage systems require high-performance networking and the ability to run multiple types of storage software for flash management, RAID or erasure coding, access control, configuration, encryption, compression/deduplication, monitoring, etc. The clients or hosts connecting to storage require storage traffic acceleration over either RDMA or TCP, plus storage virtualization. Traditionally, these tasks all run on general-purpose x86 CPUs, consuming more and more CPU cores as the networking speeds get faster and security requirements become more stringent.

A recent trend is to deploy NVIDIA BlueField DPUs inside both the storage servers and the clients/hosts accessing the storage. In the hosts or clients accessing the storage, BlueField offloads and accelerates the popular NVMe-oF storage protocol and can also offload encryption of data in motion. The BlueField SNAP technology virtualizes network storage, making it appear as a high-performance local flash device. These functions free up CPU cycles on the host, allowing more CPU cores to run applications instead of networking and storage virtualization tasks.

Both monolithic and scale-out storage often have a group of front-end controllers managing access to back-end storage enclosures. In front-end storage controllers, BlueField replaces the network cards and offloads data movement, encryption, and RDMA tasks from the CPU, allowing the controllers to either support more users or run more storage software. In dedicated storage enclosure nodes, BlueField can manage all storage software and control functions, replacing the x86 CPU and network cards, resulting in significant savings to size, weight, and power (SWaP) requirements.

Examples of BlueField Acceleration for Storage

Many newer storage systems and hyperconverged infrastructure (HCI) solutions are using the BlueField DPU to manage network traffic, RDMA, encryption, and storage management software.

- **DDN Infinia** offers software-defined, scale-out enterprise storage optimized for multi-tenancy, containerization, strong security, and simplified management. Infinia can support the BlueField DPU to offload and accelerate networking tasks, resulting in higher storage performance and efficiency.
- Hitachi Vantara has begun incorporating the BlueField DPU into its **VSP E Series** arrays with plans to expand functionality over time.
- The **Nebulon** cloud infrastructure solution uses the BlueField DPU as the basis for its Medusa2 Services Processing Unit. Offloads to BlueField allow Medusa2 to support higher performance with more virtual machines.
- **VAST Data Platforms** use BlueField DPUs, replacing the x86 CPU and resulting in a 50 percent savings in size and 30 percent savings in power consumption, while still doubling network connectivity speed from the previous generation.
- WEKA storage is adding BlueField DPU support. The WEKA client container is able to run on BlueField so it doesn't consume server CPU resources. Learn more about this work here: Weka.io on BlueField-3 (**Tech Preview: WEKA on NVIDIA BlueField-3 DPUs – WEKA**)

Ready to Get Started?

For more information on BlueField-accelerated storage solutions, please contact our partners listed above or your NVIDIA sales representative.

To learn more about the NVIDIA BlueField Networking Platform, visit **nvidia.com/dpu**.

NVIDIA